



Lenovo ThinkAgile CP Planning Guide



Models: CP 4000, CP 6000

Note

Before using this information and the product it supports, be sure to read and understand the safety information and the safety instructions, which are available at the following address:

http://thinksystem.lenovofiles.com/help/topic/safety_documentation/pdf_files.html

In addition, be sure that you are familiar with the terms and conditions of the Lenovo warranty for your solution, which can be found at the following address:

<http://datacentersupport.lenovo.com/warrantylookup>

Second Edition (March 2020)

© Copyright Lenovo 2018, 2020.

LIMITED AND RESTRICTED RIGHTS NOTICE: If data or software is delivered pursuant to a General Services Administration "GSA" contract, use, reproduction, or disclosure is subject to restrictions set forth in Contract No. GS-35F-05925.

Contents

Contents	i	Networking system limits	14
Chapter 1. ThinkAgile CP Overview	1	Power considerations	14
ThinkAgile CP4000 series overview	3	Storage considerations	15
ThinkAgile CP6000 series overview	6	Networking considerations	15
ThinkAgile CP documentation and related content	9	Pre-install network configuration worksheet	17
Chapter 2. Planning for ThinkAgile CP	11	Firewall Settings	26
Requesting an organization-specific certificate	11	Licensing considerations	26
System limits	13	Managing the cloud	27
Compute system limits	13	Updating ThinkAgile CP firmware	27
Storage system limits	13	Over the Air Updates	27
		Chapter 3. Engaging Lenovo Services	29

Chapter 1. ThinkAgile CP Overview

This product overview provides technical information on the Lenovo ThinkAgile CP Series cloud platform, its key features and specifications, components and options, and configuration guidelines.

Lenovo ThinkAgile CP Series is an all-in-one, composable cloud with an integrated application marketplace (Lenovo Cloud Marketplace) and end-to-end automation, delivering a turn-key cloud experience in your own data center. The ThinkAgile CP Series uses modular compute, storage, and networking components paired with the cloud virtualization software to create pools of IT resources, independently scaling and allocating capacity, and automatically configuring resources to fulfill application requirements.

Due to its software-defined modular architecture, the ThinkAgile CP Series platform can be scaled easily by adding more compute and storage resources independently of each other as your needs grow. Suggested workloads for the ThinkAgile CP Series include web services, virtual desktop infrastructure (VDI), enterprise applications, OLTP and OLAP databases, data analytics, application development, cost-optimized virtualization, containers, and other back-office applications.

Key Features

Lenovo ThinkAgile CP Series offers the following key features:

- Modular and scalable ThinkAgile CP Series configurations of an on-premises cloud platform featuring the second generation of the Intel Xeon Processor Scalable Family that are designed to optimize your workload's performance and provide the IT agility for your business demands.
- Factory-preloaded ready-to-go cloud platform that is delivered with all the infrastructure you need for your workloads: Modular and scalable physical and virtual compute, network, and storage resources, and the cloud software.
- Lenovo deployment services that are included with the platform help get customers up and running quickly.
- The cloud software provides a scalable software-defined infrastructure (SDI) that simplifies cloud deployments with integrated compute, storage, networking, security, and management services that manage application infrastructure and automate and orchestrate workload provisioning.
- Extensive security features, such as data at rest encryption, virtualized network and VM-level firewalls, and two-factor authentication help customers meet the most stringent security requirements.
- Centralized cloud-based management automates discovery, deployment, and configuration of cloud resources, and provides automated, non-disruptive software updates.
- Lenovo ThinkAgile Advantage provides a single point of contact for all support issues and integrates support chat and virtual technical assistance into your management interface.

Components

The ThinkAgile CP Series consists of the following components:

- **Storage Blocks**

A Storage Block is a 2U storage enclosure with up to 24 PCIe NVMe SSDs and two controllers for high availability and redundancy. A storage block runs the cloud software that delivers data storage layer to the cloud. It simplifies storage management, provides enterprise-class storage functionality, and enables linear performance and capacity scaling.



Figure 1. Storage block

- **Compute Blocks**

A Compute Block is a modular 2U enclosure that contains up to four nodes, and it delivers processor and memory resources to the cloud. The compute nodes run the cloud hypervisor that combines open KVM-based virtualization software, hardware integration, and automation to orchestrate and deliver an end-to-end compute platform.




Figure 2. Compute block

- **CP Interconnect**

The Interconnect centralizes connectivity of your on-premises infrastructure to the Cloud Controller and acts as the entry point into your existing network, which connects the Cloud Controller to your environment. The CP Interconnect consists of one or two high-density, ultra-low-latency 10/40 GbE network devices and use specialized cloud software which provides the fabric for your composable environment, automating discovery, onboarding, and device management. Integrated network virtualization secures virtual data center and applications.




Figure 3. CP Interconnect

- **Management Switch**

The 1 GbE management switch is required to provide out-of-band (OOB) management for the Interconnect, Compute Blocks, and Storage Blocks. The switch can be supplied by Lenovo (ThinkSystem NE0152T RackSwitch), or customers may provide their own switch.



Figure 4. Management switch

- **Cloud Controller**

Cloud Controller orchestrates and manages the on-premises infrastructure and workloads; however, it resides in the cloud. It automates and orchestrates infrastructure provisioning in real time to achieve complete and secure vertical hardware, software and workload integration. It federates and abstracts all physical hardware into a private cloud service.

Cloud Controller provides a single point of management across an unlimited number of ThinkAgile CP stacks with role-based access control, two-factor authentication, and secure HTML and RESTful API interfaces.

ThinkAgile CP4000 series overview


The Lenovo ThinkAgile CP4000 series scales as an initial data-management footprint for managed service provider and cloud service provider environments that start on a smaller scale but are poised for growth. The CP4000 series also supports business continuity when deployed remotely as an off-site backup for disaster recovery.

The Lenovo ThinkAgile CP4000 series consists of the following components:

- [“Storage block” on page 4](#)
- [“Compute block” on page 4](#)

- “Interconnect switch” on page 5
- “Management switch” on page 5
- “Services” on page 5
- “Software” on page 5

Storage block




- 1 Storage Block (SB-S10)
- Dual controller - single CPU per controller
- 8x SFP Base-T SFP+ transceivers
- 8x CAT6 cables, 1.5m and other lengths permitted
- 2x CAT5e cables, 1.5m and other lengths permitted
- 32 GB RAM per controller
- 8, 16, 24 drives (8 drives per drive pack)
- 800 GB, 1.6 TB NVMe drives per controller
- Drives must be consistent within a storage block
- Dimensions: 87.6 x 446.4 x 857.2 mm (3.4 x 17.6 x 33.7 inches)
- Maximum weight: 42.0 kg (92.6 lb)

Compute block



- 1 D2 enclosure, fully populated PSU, fan, 220V with 10Gb SFP+ network module
- Minimum 2 SD530 nodes, maximum 4 nodes
- 128 GB, 256 GB RAM per node
- 8+ cores per CPU, 1-2 sockets filled AXQT, and others
- 2x SFP+ DAC cables per node, 1.5m and other lengths permitted
- Boot: 1x SATA 240 GB
- Dimensions: 41.0 x 562.0 x 222.0 mm (1.7 x 22.2 x 8.8 inches)
- Maximum weight: 55 kg (121.2 lb)

Interconnect switch



- 1 or 2 CP-I-10 interconnect switches
- 48x 10 Gb SFP+ fabric ports
- 6x 40 Gb QSFP+ fabric ports
- Dimensions: 438.4 x 473 x 43.4 mm (17.26 x 18.62 x 1.71 inches)
- Maximum weight: 8.95 kg (19.73 lb), with two installed PSUs

Management switch

Lenovo ThinkSystem NE0152T RackSwitch



Note: A dedicated out-of-band (OOB) management switch is required for this solution. If the Lenovo NE0152T is not used, you must provide an OOB management switch with the correct configuration.

Services

Lenovo Deployment Services are mandatory

Lenovo Hardware Installation Service is optional

Software

Software is preloaded on compute, storage, and switch.

ThinkAgile CP6000 series overview

The Lenovo ThinkAgile CP6000 series offering can be scaled up to 40 compute nodes to support mainstream enterprise workloads and mixed-workload environments.

The Lenovo ThinkAgile CP6000 series consists of the following components:


- “Storage block” on page 6
- “Compute block” on page 7
- “Interconnect switch” on page 7
- “Management switch” on page 7
- “Services” on page 8
- “Software” on page 8

Storage block




- 1 Storage Block - a maximum 5 blocks (SB-D20)
- Dual controller - 2CPU per controller
- 8x SFP Base-T SFP+ adapters
- 8x CAT6 cables per storage block, 1.5m and other lengths permitted
- 64 GB RAM per controller
- 8, 16, 24 drives (8 drives per drive pack)
- 1.6 TB, 3.2 TB, 6.4 TB NVMe drives per controller
- Drives must be consistent within a controller
- Dimensions: 87.6 x 446.4 x 857.2 mm (3.4 x 17.6 x 33.7 inches)
- Maximum weight: 42.0 kg (92.6 lb)

Compute block



- 1 D2 enclosure, fully populated PSU, fan, 220V with 10Gb SFP+ network module — maximum 10 enclosure
- Minimum 2 SD530 nodes, default 4, maximum 40 nodes
- 128 GB, 256 GB, 384 GB, 512 GB, 768 GB, 1024 GB RAM
- 8+ cores per CPU, 1-2 sockets filled AXQT, and others
- 2x SFP+ DAC cables per node, 1.5m and other lengths permitted
- Boot: 1x SATA 240 GB
- Dimensions: 41.0 x 562.0 x 222.0 mm (1.7 x 22.2 x 8.8 inches)
- Maximum weight: 55 kg (121.2 lbs)

Interconnect switch



- 2 CP-I-10 Interconnect switches
- 48x 10 Gb SFP+ fabric ports
- 6x 40 Gb QSFP+ fabric ports
- Dimensions: 438.4 x 473 x 43.4 mm (17.26 x 18.62 x 1.71 inches)
- Maximum weight: 8.95 kg (19.73 lb), with two installed PSUs

Management switch

Lenovo ThinkSystem NE0152T RackSwitch



Note: A dedicated out-of-band (OOB) management switch is required for this solution. If the Lenovo NE0152T is not used, you must provide an OOB management switch with the correct configuration.

Services

Lenovo Deployment Services are mandatory

Lenovo Hardware Installation Service is optional

Software

Software is preloaded on compute, storage, and switch.

ThinkAgile CP documentation and related content

Information about the ThinkAgile CP product is available online and in downloadable Portable Document Format (PDF) files. In addition, documentation is available for each of the components in the product.

ThinkAgile CP publications

The most recent ThinkAgile CP documentation is at the following website:

http://thinkagile.lenovofiles.com/help/topic/thinkagile_cp/thinkagile_cp_introduction.html

The following ThinkAgile CP product publications are available in PDF format. The entire content of these publications is contained in the web-based documentation set.

- **Hardware Installation Guide**

Describes procedures for installing and cabling Lenovo ThinkAgile CP Series components.

- **Part Replacement and Component Maintenance Procedures**

Describes procedures for removing and replacing Lenovo ThinkAgile CP Series components.

ThinkAgile CP product page

For information about the current ThinkAgile CP models, refer to the [Lenovo ThinkAgile CP product page](#). See the following website:

<https://www.lenovo.com/us/en/data-center/software-defined-infrastructure/ThinkAgile-CP-Series/p/WMD00000362>

ThinkAgile CP Support

For the latest information and links to helpful content, refer to the [Lenovo Support website](#). You can search for support information for ThinkAgile CP and other Lenovo Data Center products. See the following website for information related to ThinkAgile CP:

- CP 4000:

<https://datacentersupport.lenovo.com/us/en/products/solutions-and-software/thinkagile-cp/thinkagile-cp4000/7y34>

- CP 6000:

<https://datacentersupport.lenovo.com/us/en/products/solutions-and-software/thinkagile-cp/thinkagile-cp6000/7y34>

Chapter 2. Planning for ThinkAgile CP

The following considerations and limitations apply to the ThinkAgile CP solution.

Firmware and “Best Recipe” adherence

Lenovo maintains a ThinkAgile CP firmware “Best Recipe”, which identifies the supported levels for the various components. Contact Lenovo Support for more details.

Certificate Request

Attention: Since it can take up to two (2) business days for Lenovo to sign the GNU Privacy Guard (GPG) keys used in the ThinkAgile CP solution certification, we recommend that you first complete steps 1 through 5 from [“Requesting an organization-specific certificate” on page 11](#).

Requesting an organization-specific certificate

For new customers, a certificate must be requested at least two (2) business days in advance by sending a request to certificates@tacp.zendesk.com.

The following information must be provided in the e-mail:

Required details	Example
Requester Name	<i>John Doe</i>
Requester E-mail	<i>john.doe@acme.com</i>
Date	<i>2018-12-31</i>
Company Name	<i>ACME Corporation</i>
Organization Name	<i>MSP Stacks</i>
GPG Public key	<i>Public key name (publicKey.asc)</i>

Note: An organization provides a way to organize and manage your ThinkAgile CP hardware resources. The organization identifier is used to identify the hardware stack in the ThinkAgile CP Cloud Controller. Each organization identifier must be unique. For more information about organizations within the ThinkAgile CP Cloud Controller, see [Manage Organizations](#).

The Lenovo team will reply with the certificate and the instructions to install.

To create and use GNU Privacy Guard (GPG) keys, follow these steps:

Step 1. Execute the following command to create a new GPG key:

```
gpg2 --gen-key
```

Alternatively, you can execute the following command:

```
gpg2 --gen-key
```

Note: This prompts you for an interactive setup. The following example shows sample output and answers for reference.

```
gpg2 --gen-key
```

gpg (GnuPG) 2.0.22; Copyright (C) 2013 Free Software Foundation, Inc.
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.

```
gpg: directory `/root/.gnupg' created
gpg: new configuration file `/root/.gnupg/gpg.conf' created
gpg: WARNING: options in `/root/.gnupg/gpg.conf' are not yet active during this run
gpg: keyring `/root/.gnupg/secring.gpg' created
gpg: keyring `/root/.gnupg/pubring.gpg' created
Please select what kind of key you want:
  (1) RSA and RSA (default)
  (2) DSA and Elgamal
  (3) DSA (sign only)
  (4) RSA (sign only)
Your selection? 1
RSA keys may be between 1024 and 4096 bits long.
What keysize do you want? (2048) 4096
Requested keysize is 4096 bits
Please specify how long the key should be valid.

    0 = key does not expire
    <n> = key expires in n days
    <n>w = key expires in n weeks
    <n>m = key expires in n months
    <n>y = key expires in n years
Key is valid for? (0) 2y
Key expires at Sat Aug  8 14:33:52 2020 UTC
Is this correct? (y/N) y
```

GnuPG needs to construct a user ID to identify your key.

```
Real name: John Doe ( ACME )
Email address: jdoe@acme.com
Comment: ACME GPG
You selected this USER-ID:
  "John Doe ( ACME ) (ACME GPG)<jdoe@ACME.com>"
Change (N)ame, (C)omment, (E)mail or (O)kay/(Q)uit? 0
You need a Passphrase to protect your secret key.
```

Step 2. Locate your public key IDs:

`gpg2 --list-public-keys <requester e-mail>`, where the `<requester e-mail>` is the e-mail address used to request the certificate from `certificates@tacp.zendesk.com`.

The following is an example showing the command output:

```
gpg2 --list-public-keys john.doe@acme.com

pub  4096R/ECA7E403 2018-08-08 [expires: 2020-08-07]
uid                               John Doe ( ACME Inc ) <john.doe@acme.com>
sub  ....
```

Note: In the above output, the key is **ECA7E403**.

Step 3. Export the key and send your keys to the key server:

```
gpg2 --send-keys --keyserver pool.sks-keyservers.net <GPG key ID>
```

For example:

```
gpg2 --send-keys --keyserver pool.sks-keyservers.net ECA7E403

gpg: sending key C264BE84ECA7E403 to hkp://pool.sks-keyservers.net
```

Step 4. Extract the GPG key:


```
gpg --output publicKey.asc -- armor -- export <requester e-mail>, where the <requester e-mail> is the e-mail address used to request the certificate from certificates@tacp.zendesk.com.
```

Note: Use **publicKey.asc** as the name of the public key.

Step 5. Request the new certificate from *certificates@tacp.zendesk.com* as described at the beginning of this section.

Step 6. Please wait for the public key to be signed by Lenovo.

Note: It can take up to two (2) business days to sign the GPG key.

Step 7. Decrypt the signed keys obtained from Lenovo. The result of the decryption process is a .TAR file.

Note: This must be done in the same directory where you extracted the key.

```
gpg2 --output <New-TAR-File-Name> -- decrypt <Signed-File-Obtained-From-Lenovo>
```

For example:

```
gpg2 --output lenovo-thinkagile-test.controller-broker-client.pki.cp.lenovo.com-client_db.tar.gz --decrypt lenovo-thinkagile-test.controller-broker-client.pki.cp.lenovo.com-client_db.tar.asc
```

System limits

Use this information to determine system limits within the ThinkAgile CP solution.

Compute system limits

Use this information to determine system limits for a ThinkAgile CP compute block.

Component	Limit description
Number of vNICs	64 vNICs per VM, or application instance
Number of application instances per hypervisor	No theoretical limit on the number of VMs. This does rely on the number of cores.

Storage system limits

Use this information to determine system limits for a ThinkAgile CP storage block.

All recommended limits listed below are per Storage Block. Adding multiple Storage Blocks within a federated storage pool will generally increase these limits linearly, excluding VM specific limits such as VM disk limit. These recommended limits are based on system stability, hard-coded system limits, as well as time limits to complete each single action (for example, backup/disk creation).

Component	Limit description
Inactive blocks, disks, backups, or base images (per storage block)	1 million
Active storage disks exported to VMs (per storage block)	10,000
Disks per VM	200

Component	Limit description
Backups per VM	Limited to the total number of blocks (1 million) and storage block capacity per storage block
Disaster recovery backups per VM	Limited to the total number of blocks (1 million) and storage block capacity per storage block

Networking system limits

Use this information to determine system limits for a ThinkAgile CP interconnect switch.

Component	Limit description
Number of VNETs that can be created	There is no limit to the number of VNETs that you can create for an Interconnect, as long as you have resources to create an NFV instance. Note: VNETs are represented by a 64-bit unsigned integer; theoretically, there can be 18 quintillion VNETs.
Number of VLANs that the interconnect switch supports	The interconnect switch can support 4,096 VLANs. However, there may be additional limitations based on the STP version you are running. The ThinkAgile CP interconnect switch uses five VLANs internally: <ul style="list-style-type: none"> • MLAG peerlink Link VLAN (ports 49 and 50) • Hardware management network VLAN (port 47) • Uplink VLAN (port 48, but varies per customer basis) • Hypervisor VLAN (also used by the SAN) (port 1-44) • Switch out-of-band
STP	An interconnect switch can support 64 instances of PVST+ (59 customer VLANs, excluding the five reserved internal VLANs described above). MSTP can support 16 instances.
vNICs	You can create up to 64 NICs and have them all be in different VLANs.

Power considerations

Power requirements and issues to consider when planning for your Lenovo ThinkAgile CP installation.

In planning for the power needs of the ThinkAgile CP solution, consider the following issues:

- Ensure that the correct number of power receptacles is available based on the quantity of nodes.
- Ensure that the power cords ordered match the power receptacles in the customer's rack.
- Ensure that power cords are ordered with lengths that allow them to connect to the customer-provided PDUs.

Table 1. Power specifications for each powered component type in stack

Device	Number of power supplies	PSU voltage	PSU max voltage	Heat generation (BTUs)
Network interconnect CP-I-10	2	100–240 VAC	400 watts	
Compute block CP-CB-10	2	220V	2000 watts	<p>Minimum configuration (with one minimal configuration node): 604.1 BTU per hour (177 watts)</p> <p>Maximum configuration (with four maximal configuration nodes): 7564.4 BTU per hour (2610 watts)</p>
Storage block CP-SB-10	2	220V	2000 watts	<p>Minimum configuration (with one minimal configuration node): 604.1 BTU per hour (177 watts)</p> <p>Maximum configuration (with four maximal configuration nodes): 7564.4 BTU per hour (2610 watts)</p>

Storage considerations

Storage issues to consider when planning for your Lenovo ThinkAgile CP installation.

In planning for the storage needs of the ThinkAgile CP solution, consider the following issues.

Choose the storage blocks

Choose the quantity and configuration of storage blocks. For each storage block, choose:

- The size of the storage drive.
- The quantity of storage drive packs, which come in quantities of 8 drives.

Notes:

- All drives within a storage block must have the same capacity.
- Storage block capacity can be increased after initial purchase by adding more 8-drive packs of the same capacity.

Networking considerations

Networking issues to consider when planning for your Lenovo ThinkAgile CP installation.

You need to consider the following topics before proceeding with the installation of the ThinkAgile CP solution.

System network cables

ThinkAgile CP delivers with network cables having lengths to enable installation of all equipment in the stack into a single customer rack.

Data Center Uplink requirements

- Cables - ThinkAgile CP can be ordered with optional cables for uplinks to the customer's data center network. If these cables are insufficient, plan the correct cable types and lengths with the customer.
- Adapters - ThinkAgile CP can be ordered with optional adapters to provide an uplink to the customer's data center through the solution interconnects. You must plan and purchase additional adapters for the customer's existing data center network hardware.
- Routes - prior to installation, the customer must decide on the routing protocol used to connect to the data center network, and must configure the routes on the hardware connected to the ThinkAgile CP interconnect.

Pre-install network configuration

The following figure provides an overview of the network planning to be performed before the installation of ThinkAgile CP.

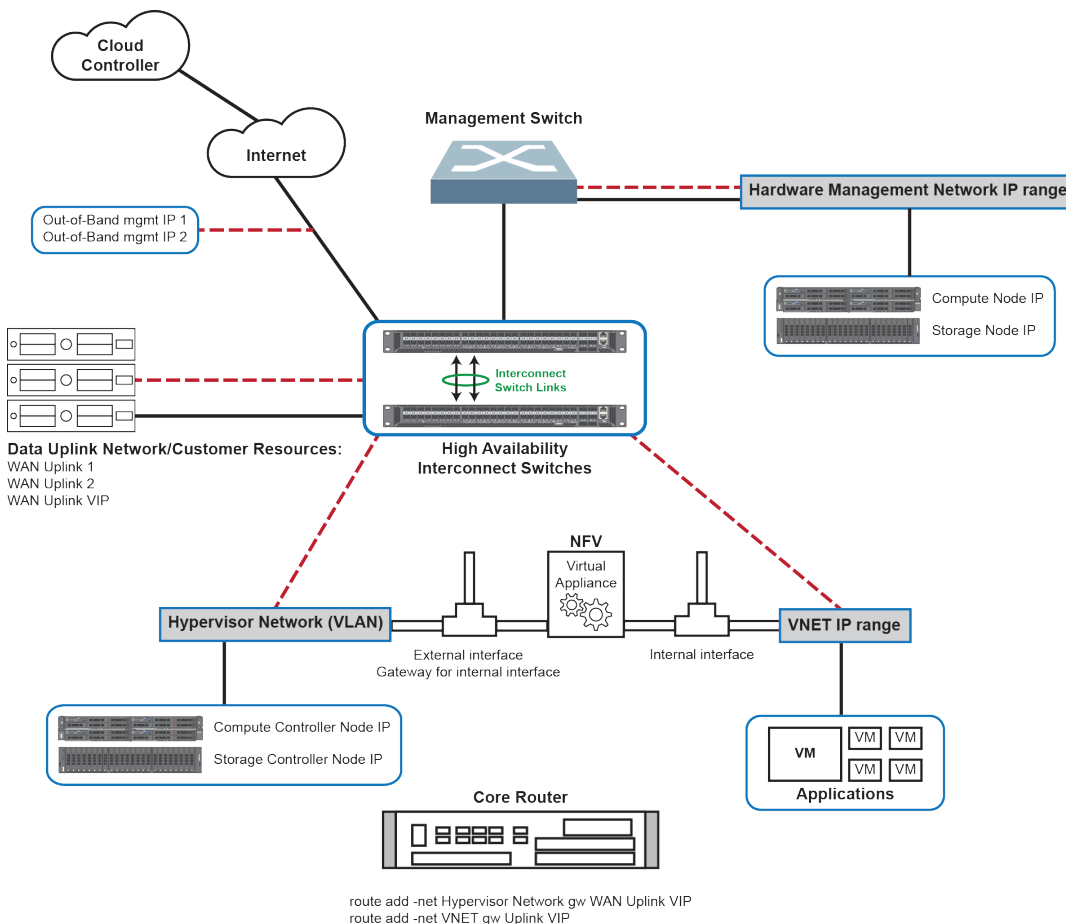


Figure 5. Network planning diagram

VLANs

- Prior to installation, the customer must provide a list of VLANs to trunk into the ThinkAgile CP network.
- We recommend trunking VLANs into the ThinkAgile CP network during the initial setup process.

IP address ranges

- Prior to installation, the customer must provide at least one IP address range for use by the ThinkAgile CP system to automate IP address allocation.
- We recommend assigning one or more dedicated class B or class C networks to ThinkAgile CP.

Firewall rules

- Prior to installation, the customer must configure their network firewall settings to allow the ThinkAgile CP system access to the internet. This is required for some functions to properly work.
- We recommend configuring both outgoing and incoming firewall rules.

Pre-install network configuration worksheet

Before installing the ThinkAgile CP solution, you need to plan your deployment by completing a pre-install network configuration worksheet.

Physical interfaces and cabling types

When planning your ThinkAgile CP deployment, consider the following:

Data Center Uplink Port Type	10 GB or 40 GB
Cable Type	SFP+ fiber ports (10 GbE) or QSFP+ fiber ports (40 GbE)
OOB Interconnect Port Type	RJ45
Data Center PDU power Connector Type	C14

Worksheet considerations

This section covers the considerations required for filling out the pre-install network configuration worksheet.

Take into account the following when filling out the worksheet:

- **Data Center Uplink**

The data center uplink is the main data path of all northbound and southbound traffic in support of the ThinkAgile CP environment. This VLAN is customer defined and lives on the customer network. This network needs connectivity to the Internet (see [“Firewall Settings” on page 26](#)). It also needs a unique IP address and subnet across the ThinkAgile CP environment.

- For dual interconnect configurations, you must provide one IP address for each interconnect switch and one IP address for the Virtual Router Redundancy Protocol (VRRP).
- For single interconnect configurations, only one IP address is required.

- **OOB Interconnect network**

This network is used to manage the interconnect switches; it provides hardware health reporting and management between the interconnect switches and the cloud controller. This network needs connectivity to the Internet (see [“Firewall Settings” on page 26](#)). It also needs a unique IP address and netmask across the ThinkAgile CP environment.

- **Hypervisor network**

The hypervisor network is used to provide internal connectivity among the host operating systems of the compute and storage controllers, and the virtual networking (VNET) NFV appliances. This network must be unique and routable from the customer network.

- **Hardware management network**

This network connects the System Management Module (SMM) for each compute enclosure, the Cloud Controller for each compute node, and the Baseboard Management Controller (BMC) for each storage controller.

- **Multi-enclosure link aggregation network (MLAG)**

This network operates on the Interconnect Switch Link (ISL) between the two interconnects and is used for internal control communication associated to the multi-enclosure link aggregation protocol between the interconnect switches. This network only requires two IP addresses, and by default, the system uses addresses taken from a reserved range.

- **Variable names**

Variable names are shown in italics in worksheet tables. Those names correspond to parameters that you need to specify when deploying the ThinkAgile CP solution (using the *ThinkAgile CP Deployment Guide*).

Prior to installation, use the following table as a guide to the configuration details required for a successful ThinkAgile CP deployment.

Table 2. Configuration details guidelines

Network	IP Address Requirements	Assignment Details	Address Type
Data Center Uplink	1 IP address per Interconnect 1 Virtual IP Address (VRRP)	Statically configured at installation time	public ¹
OOB Interconnect	1 IP address per Interconnect	Statically configured at installation time	public
Hardware management network VLAN ²	1 IP address per Interconnect	Statically configured at installation time	private
Hardware management network IP Range ²	1 IP address per compute node's XCC 1 IP address per compute block SMM 1 IP address per storage controller	Dynamically configured by the DHCP server on the interconnects	private
Hypervisor VLAN	1 IP address per Interconnect 1 Virtual IP Address (VRRP)	Statically configured at installation time	public

1. The term *public* means routable in the customer data center network. In other words, any host connected to this network must be able to reach these addresses.
 2. These IP addresses are public in development environments and labs, as to allow access for debugging purposes and operating system resets during development. However, for customer installations, the recommendation is to use private IP addresses, which must not be accessible.

Table 2. Configuration details guidelines (continued)

Hypervisor IP Range	1 IP address per compute node 4 IP addresses per storage controller Remaining IP addresses that can be assigned to application VMs vNICs and outside interfaces of NFV VMs only when the corresponding vNICs are configured in VLAN mode and assigned to the default VLAN	Dynamically configured by the DHCP server on the interconnects	public
MLAG Peers	1 IP address per Interconnect	Hard-coded values, which may be optionally be statically configured at installation time	private
VNET IP Ranges	Variable number of IP addresses, which are assigned application VM's vNICs configured in VNET mode	Dynamically configured by the DHCP server on the NFV VM	public or private
VNET Outside Interface (optional)	1 IP address per NFV, assigned to the outside Interface of an NFV VM configured in VLAN mode and assigned a non-default VLAN	Dynamically configured by a DHCP server on the customer data center network or statically at VNET creation time	public
Application VM's vNICs (optional)	1 IP address per vNIC, assigned to an application VM's vNIC configured in VLAN mode and assigned a non-default VLAN	Dynamically configured by a DHCP server on the customer data center network	public

VLAN conflict check

This section covers how to check for VLAN conflicts.

You need to verify that the VLANs that are currently in use are not duplicated.

Table 3. VLAN conflict check

Name	VLAN ID (Tag)
Native VLAN ID	1
External ThinkAgile CP VLANs Visible at Customer Network	<i>uplink_vlan</i> (tagged at the data uplink ports) <i>oob_vlan</i> (tagged at the out-of-band management ports)
Internal ThinkAgile CP VLANs Not Visible at the Customer Network	<i>hyp_vlan</i> (default/recommended: 4091) <i>bmc_vlan</i> (default/recommended: 4092) <i>mlag_vlan</i> (default/recommended: 4093)

Protocol interoperability

This section covers the protocol interoperability settings for ThinkAgile CP and the customer network.

The following protocol configuration options must be implemented on both the ThinkAgile CP stack and the customer network to ensure the interoperability between the two.

Table 4. Protocol interoperability

Protocol	Protocol options
Spanning Tree Protocol	None (Default), PVST+, MSTP, or RSTP
Link Aggregation Protocol	LACP, MLAG
Routing Protocol	Static or BGP

Note: It is important to determine whether the customer edge switch supports MLAG (multi-enclosure link aggregation) because this may allow for better redundancy by connecting ThinkAgile CP to two customer edge switches.

Interconnect switch configuration considerations

This section covers the interconnect switch settings to consider during the ThinkAgile CP solution planning.

Determine the following settings before installing the interconnect switch in the ThinkAgile CP stack.

Table 5. Interconnect switch configuration considerations

Setting name	Setting options
DNS Name Server	Primary DNS Name Server: 8.8.8.8 (<i>dns_primary_nameserver</i>) Secondary DNS Name Server: 8.8.4.4 (<i>dns_secondary_nameserver</i>)
NTP Server	Primary NTP Server: pool.ntp.org (<i>primary_ntp_server</i>) Secondary NTP Server: time.nist.gov (<i>secondary_ntp_server</i>)

Table 5. Interconnect switch configuration considerations (continued)

Host Name	<p>Primary Host Name: <i>hostname_primary</i></p> <p>Secondary Host Name: <i>hostname_secondary</i></p>
Interconnect Topology Configuration	<p>For more information about the interconnect configuration topologies, see the sections related to connecting to a customer data center network in the <i>ThinkAgile CP Hardware Installation Guide</i>. The following customer interconnect configuration types are supported:</p> <ul style="list-style-type: none"> • Single interconnect (single-direct link, dual-direct link, or LAG direct link) • Dual interconnect (X-Link Option 1, X-Link Option 2A, or X-Link Option 2B) <p><i>INTERCONNECT_CONFIG</i> variable:</p> <ul style="list-style-type: none"> • Option1 (create single MLAG across all uplink ports) • Option2 (create two MLAG ports out of all uplink ports) <p>Specify the value for the <i>INTERCONNECT_CONFIG</i> variable in the <i>ThinkAgile CP Deployment Guide</i>.</p> <p>Note: For a single Interconnect, you must select <i>option1</i>.</p>

Static route configuration in the customer network - Edge and core routers

This section covers the static route configuration considerations for edge and core routers in the customer network.

The following configuration is required if no dynamic routing protocols (e.g. BGP) are enabled between the stack's interconnects and the customer network.

Table 6. Static route configuration in the customer network - edge and core routers

Destination Subnet	Netmask	Gateway	Description
<i>hyp_range_subnet</i>	<i>hyp_range_netmask</i>	<i>wan_vip</i>	Route to reach the hypervisor network, which is used for management and control communication between the ThinkAgile CP stack and the Cloud Controller.
<i>vnet_range_subnet1</i> ... <i>vnet_range_subnetN</i> (optional)	<i>vnet_range_netmask1</i> ... <i>vnet_range_netmaskN</i> (optional)	<i>wan_vip</i>	One route per VNET, which must be reachable from the customer data center network. ³

3. These routes are only required under two conditions: no dynamic routing protocol is enabled between the ThinkAgile CP stack and the customer data center network; and you want to allow applications running in these VNETs to access hosts in the customer data center network. In other words, the VNET into consideration is using a public IP address range.

IP address assignment – customer facing networks

This section covers the guidelines when assigning IP addresses to customer facing networks.

Take into account the following when filling out the worksheet table in this section:

- **Data Center Uplink**

The data center uplink is the main data path of all northbound and southbound traffic in support of the ThinkAgile CP environment. This VLAN is customer defined and lives on the customer network. This network needs connectivity to the Internet (see “[Firewall Settings](#)” on page 26). It also needs a unique IP address and subnet across the ThinkAgile CP environment.

- For dual interconnect configurations, you must provide one IP address for each interconnect switch and one IP address for the Virtual Router Redundancy Protocol (VRRP).
- For single interconnect configurations, only one IP address is required.

- **OOB Interconnect network**

This network is used to manage the interconnect switches; it provides hardware health reporting and management between the interconnect switches and the cloud controller. This network needs connectivity to the Internet (see “[Firewall Settings](#)” on page 26). It also needs a unique IP address and netmask across the ThinkAgile CP environment.

- **Variable names**

Variable names are shown in italics in worksheet tables. Those names correspond to parameters that you need to specify when deploying the ThinkAgile CP solution (using the *ThinkAgile CP Deployment Guide*).

- **Network masks**

In worksheet tables, network masks must be entered in the Netmask or Prefix column using dotted decimal format X.X.X.X (for example: 10.1.1.2).

- **Network prefixes**

In worksheet tables, network prefixes must be entered in the Netmask or Prefix column in prefix format /<number of bits> (for example: /23).

The following are IP address assignment considerations for customer facing networks:

Table 7. IP address assignment – customer facing networks

Name	IP Range	Netmask or Prefix	Gateway	VRRP Virtual IP (VIP)	Switch Port Mode	VLAN ID (Tag)
Data Center Uplink IP 1	<i>wan_ip1</i>	<i>wan-netmask</i>	<i>wan_gw</i>	<i>wan_vip</i>	N/A	<i>uplink_vlan</i>
Data Center Uplink IP 2	<i>wan_ip2</i>					
OOB Interconnect 1	<i>oob_ip1</i>	<i>oob_prefix</i>	<i>oob_gw</i>	N/A	Access	<i>oob_vlan</i>
OOB Interconnect 2	<i>oob_ip2</i>					

Notes:

- The variables *wan_ip1*, *wan_ip2*, *wan_vip* and *wan_gw* must belong to the same subnet.
- The variable *wan_gw* refers to the IP address assigned to the customer edge router attached to the stack’s interconnects via the data uplinks.

- The variables *oob_ip1*, *oob_ip2*, and *oob_vip* must belong to the same subnet.

IP address assignment – ThinkAgile CP internal networks

This section covers the guidelines when assigning IP addresses to ThinkAgile CP internal networks.

Take into account the following when filling out the worksheet table for ThinkAgile CP internal networks

- **Hypervisor VLAN**

This VLAN is used for internal communication within the ThinkAgile CP stack and not customer accessible. The value assigned to this VLAN is used in the *ThinkAgile CP Deployment Guide*. We recommend using the default value 4091 for this VLAN. Nevertheless, its value can be configured during installation using the variable named *hyp_vlan*.

- **Hypervisor IP Range**

This range is used to assign IP addresses to the operating systems of the compute and storage controllers. The network associated to this IP range must be unique and routable from the customer data center network. Under special circumstances, this range is used to assign IP addresses to the outside interface of the NFV VMs or application VMs as long as the corresponding vNICs are configured in VLAN mode and assigned to the default VLAN.

- **Hypervisor IP Range Gateway**

The gateway for all hosts in the hypervisor IP range is the Hypervisor VLAN Virtual IP (*hyp_vlan_vip*).

- **Hardware management network VLAN**

This VLAN is internal to ThinkAgile CP and not customer accessible.

- **Hardware management network VLAN**

We recommend using the default value 4092 for this VLAN. Nevertheless, its value can be configured during installation using the variable named *bmc_vlan*.

- **Hardware management network IP Range**

This IP range is used for internal communication within the ThinkAgile CP stack and not customer accessible. It is used to assign addresses to the BMC devices within the compute and storage controllers.

- **Hardware management network IP Range Gateway**

The gateway for all hosts in this IP range is the hardware management network VLAN IP address (*bmc_vlan_ip*).

- **MLAG Peer Link**

We recommend using the default values 100.64.253.1 and 100.64.253 for the MLAG Peer Link endpoint addresses. These IP addresses belong to the block 100.64.0.0/10, which was assigned by Internet Assigned Numbers Authority (IANA) to be used as shared address space range, typically for carried grade NAT. In the context of ThinkAgile CP, these IP addresses are private, and the MLAG network is completely isolated from the customer data center network. Although ThinkAgile CP uses these values by default, the user has the flexibility to modify them during installation using the variables named *mlag_vlan_ip1*, *mlag_vlan_ip2* and *mlag_vlan_prefix*.

- **MLAG VLAN**

We recommend using the default value 4093 for this VLAN. Nevertheless, its value can be configured during installation using the variable named *mlag_vlan*.

- **MLAG VLAN Prefix**

We recommend using the default value 30 for this prefix. Nevertheless, its value can be configured during installation using the variable named *mlag_vlan_prefix*.

- **Variable names**

Variable names are shown in italics in worksheet tables. Those names correspond to parameters that you need to specify when deploying the ThinkAgile CP solution (using the *ThinkAgile CP Deployment Guide*).

- **Network masks**

In worksheet tables, network masks must be entered in the Netmask or Prefix column using dotted decimal format X.X.X.X (for example: 10.1.1.2).

- **Network prefixes**

In worksheet tables, network prefixes must be entered in the Netmask or Prefix column in prefix format /<number of bits> (for example: /23).

The following are IP address assignment considerations for ThinkAgile CP internal networks:

Table 8. IP address assignment – ThinkAgile CP internal networks

Name	IP Range	Netmask or Prefix	Gateway	VRRP Virtual IP (VIP)	VLAN ID (Tag)
Hypervisor VLAN	<i>hyp_vlan_ip</i>	<i>hyp_vlan_prefix</i>	N/A	<i>hyp_vlan_vip</i>	
Hypervisor IP Range	<i>hyp_range_subnet</i> <i>hyp_range_start</i> <i>hyp_range_end</i>	<i>hyp_range_netmask</i>	<i>hyp_vlan_vip</i>	N/A	<i>hyp_vlan</i> Default: 4091 (optional)
Hardware management network VLAN	<i>bcm_vlan_ip</i>	<i>bmc_vlan_prefix</i>	N/A		
Hardware management network IP Range	<i>bmc_range_subnet</i> <i>bmc_range_start</i> <i>bmc_range_end</i>	<i>bmc_range_netmask</i>	<i>bmc_vlan_ip</i>	N/A	<i>bmc_vlan</i> Default: 4092 (optional)
MLAG Peer Link	<i>mlag_vlan_ip1</i> <i>mlag_vlan_ip2</i> Default: 100.64.253.1 and 100.64.253.2 (optional)	<i>mlag_vlan_prefix</i> Default: 30 (optional)	N/A	N/A	<i>mlag_vlan</i> Default: 4093 (optional)

IP address assignment – VNETs

This section covers the considerations

Take into account the following when filling out the worksheet table for VNET IP address assignment:

- **VNET IP Ranges**

For VNET that you plan to configure, it is recommended that you record the following parameters, which will be required for configuring each VNET: *vnet_range_start*, *vnet_range_end*, *vnet_range_netmask*, and *vnet_range_gw*. It is important to stress that the IP ranges allocated to the VNETs must be completely

separate and isolated from the internal networks (e.g. hypervisor, out-of-band management, hardware management network, and MLAG).

- **VNET with Outside Interface in VLAN Mode Configured for Non-Default VLAN**

This is optional, and only applicable if you are planning to extend some of the VLANs configured in your network into the stack, and configure the outside interface of one of more VNETs in such VLANs. Note that this is only applicable if we configure the outside interface of the NFV VNET to use a non-default VLAN. If all the previous conditions are met, you recommend that you write down the set of VLANs in the variables [*vnet_outside_vlan1 - vnet_outside_vlanN*]. In addition, the outside interfaces of the VNETs must be assigned IP addresses, net masks, and default gateways (i.e. *vnet_outside_ip1 - vnet_outside_ipN*, *vnet_outside_netmask1 - vnet_outside_netmaskN*, *vnet_outside_gw1 - vnet_outside_gwN*). This assignment can be static upon VNET configuration, or dynamic by configuring DHCP servers/relays in the customer data center network to serve each corresponding VLAN (recommended).

- **Application VM with vNIC in VLAN Mode Configured for Non-Default VLAN**

This is optional, and only applicable if you are planning to extend some of the VLANs configured in your network into the stack, and configure at least one of the vNICs of an application VM in such VLANs. Note that this is only applicable if we configure the VM's vNIC to use a non-default VLAN. In this case, you must write down the set of VLANs in the variables [*vm_vlan1 - vm_vlanN*]. In addition, the VM's vNICs must be assigned IP addresses, net masks, and default gateways (i.e. *vm_ip1 - vm_ipN*, *vm_netmask1 - vm_netmaskN*, *vm_gw1 - vm_gwN*). Application VMs typically do not support static IP assignments on VLANs that extend to the customer data center network. As a result, dynamic assignment by configuring DHCP servers/relays in the customer data center network to serve each corresponding VLAN is required.

- **Variable names**

Variable names are shown in italics in worksheet tables. Those names correspond to parameters that you need to specify when deploying the ThinkAgile CP solution (using the *ThinkAgile CP Deployment Guide*).

- **Network masks**

In worksheet tables, network masks must be entered in the Netmask or Prefix column using dotted decimal format X.X.X.X (for example: 10.1.1.2).

- **Network prefixes**

In worksheet tables, network prefixes must be entered in the Netmask or Prefix column in prefix format /<number of bits> (for example: /23).

Table 9. IP address assignment – VNET

Name	IP Range	Netmask or Prefix	Gateway	VLAN ID (Tag)
VNET IP Ranges	<i>vnet_range_subnet1</i>			N/A
	<i>vnet_range_start1</i>			
	<i>vnet_range_end1</i>	<i>vnet_range_netmask1</i>	<i>vnet_range_gw1</i>	
	
	<i>vnet_range_subnetN</i>	<i>vnet_range_netmaskN</i>	<i>vnet_range_gwN</i>	
	<i>vnet_range_endN</i>			

Table 9. IP address assignment — VNET (continued)

VNET with Outside Interface in VLAN Mode Configured for Non-Default VLAN (optional)	<i>vnet_outside_ip1</i>	<i>vnet_outside_netmask1</i>	<i>vnet_outside_gw1</i>	<i>vnet_outside_vlan1</i>

	<i>vnet_outside_ipN</i>	<i>vnet_outside_netmaskN</i>	<i>vnet_outside_gwN</i>	<i>vnet_outside_vlanN</i>
Application VM with vNIC in VLAN Mode Configured for Non-Default VLAN (optional)	<i>vm_ip1</i>	<i>vm_netmask1</i>	<i>vm_gw1</i>	<i>vm_vlan1</i>

	<i>vm_ipN</i>	<i>vm_netmaskN</i>	<i>vm_gwN</i>	<i>vm_vlanN</i>

Firewall Settings

This section covers the firewall settings needed to properly configure the ThinkAgile CP solution.

Some functions of the Lenovo ThinkAgile CP require access to the Internet.

Outgoing Firewall Rules

If you have firewalls in your network, configure the following outgoing firewall rules to enable ThinkAgile CP to perform the following actions:

Table 10. Outgoing Firewall Rules

Destination	Protocol	Port	Description
repo-olympia.cp.lenovo.com	TCP	80, 443	Software Updates
endpoint-prod.us-east-1.cp.lenovo.com	TCP	443	Cloud Controller Alias
		5672, 5673	Cloud Controller Alias — Regional
s3.amazonaws.com	TCP	443	Lenovo Cloud Marketplace
tacp-proxy.cp.lenovo.com	TCP	All	Support Functionality
8.8.8.8 and 8.8.4.4	TCP/UDP	53	DNS Servers

Incoming Firewall Rules

If you have firewalls in your network, configure the following incoming firewall rules to enable ThinkAgile CP to perform the following actions:

Table 11. Incoming Firewall Rules

Destination	Protocol	Port	Description
Hypervisor Network IP Range	TCP	1024 — 49, 151	VM Console

Licensing considerations

The following table explains the details of the ThinkAgile CP Guardian edition.

	Guardian
Operating System (for compute nodes and storage controllers)	RHEL 7
Network Operating System for Interconnect Switches	Pica8 OS
At-Rest Encryption	✓
Compliance	
KMIP Compliant	✓
FIPS 140-2	✓
Common Criteria	✓
USGv6 (DOD IPv6)	✓
STIG	✓

Managing the cloud

The cloud is managed through the ThinkAgile CP software.

For information about managing the cloud through ThinkAgile CP, see:

https://thinkagile.lenovofiles.com/help/topic/thinkagile_cp/admin-guide-introduction.html

Updating ThinkAgile CP firmware

ThinkAgile CP firmware updates should only be planned or performed when advised by Lenovo Support. Contact Lenovo Support for information about firmware updates.

Over the Air Updates

Over the Air (OTA) updates provide new functionality, bug fixes and security upgrades to the operating system running on the interconnect switches and the storage and compute blocks as part of the ThinkAgile CP solution.

OTA updates are used to upgrade various firmware components of the operating system. Considering the impact of the updates on the customer's current workload, OTA updates represent a disruptive or a non-disruptive process.

Disruptive Updates

This type of updates impact the current workload by interrupting the normal operations of the ThinkAgile CP solution.

To complete the upgrade procedure, disruptive OTA updates may require the restart of the following ThinkAgile CP components:

- compute or storage nodes
- virtual machines
- ThinkAgile CP Cloud Controller

Non-disruptive Updates

In contrast with disruptive updates, non-disruptive updates do not have any negative impact on the current workload. They can be performed without the need to restart compute or storage nodes, or virtual machines.

OTA Update Procedure

The customer is informed of any available OTA updates via an e-mail or Cloud Controller notifications. The customer is provided with the following information about the OTA update:

- The timeline of the update
- The release notes, detailing the changes, enhancements or any new functionality introduced by the update
- The impact of the update: disruptive or non-disruptive
- The duration of the update
- Contact details for Lenovo Support

Generally, OTA updates are performed by Lenovo Support with minimal involvement from the customer. The customer works together with Lenovo Support to determine the best schedule and method to incorporate the OTA updates into their ThinkAgile CP solution.

To accommodate the needs of each individual customer, Lenovo Support analyzes the impact of the OTA update on the current workload. Then, alongside with the customer, they decide the best plan of action as to minimize the downtime of any valuable resources (such as compute nodes or VMs). They determine:

- The timeframe for the OTA update procedure
- The impact and mitigation plan
- The duration of the update procedure
- The actions required of the customer before, during and after the OTA update

Chapter 3. Engaging Lenovo Services

Lenovo Professional Services performs a specific set of activities to ensure successful deployment of the ThinkAgile CP products.

Note: The on-site engagement's success depends on careful planning and pre-work. See https://thinkagile.lenovofiles.com/help/topic/thinkagile_cp/planning_considerations.html for more details.

The on-site installation and configuration activities that Lenovo Professional Services performs include the following:

- Review the deployment worksheet completed during the pre-deployment engagement.
- Verify that the customer network is ready.
- Verify the as-shipped product configuration.
- Connect the top-of-rack switches to the customer network.
- Perform basic network configuration to integrate the ThinkAgile CP rack into the customer environment.
- Ensure that the product adheres to the ThinkAgile CP Best Recipe firmware profile.
- Configure the ThinkAgile CP user interface (e.g., SMTP, Pulse, containers) and license.
- Verify the product environment, to ensure a successful handover.
- Provide a knowledge transfer to the customer. This includes basic operational tasks and the support plan.

Lenovo Professional Services will provide information about the configured system, for your records.

The following items are not in the scope of the initial ThinkAgile CP engagement. Many of these can be handled via other available Professional Services offerings.

- Hardware installation services
- Server applications installation or configuration
- Advanced hypervisor configuration
- Complete administrative or operations training
- Data migration from existing storage systems and hypervisor environments
- Installation assistance for nonstandard hardware components or components not purchased from Lenovo. Be aware that this activity can affect support for the product.

Lenovo